

Catalogue no. 12-001-X
ISSN 1492-0921

Survey Methodology

Design-based conformal prediction

by Jerzy Wiecek

Release date: January 3, 2024



Statistics
Canada

Statistique
Canada

Canada

How to obtain more information

For information about this product or the wide range of services and data available from Statistics Canada, visit our website, www.statcan.gc.ca.

You can also contact us by

Email at infostats@statcan.gc.ca

Telephone, from Monday to Friday, 8:30 a.m. to 4:30 p.m., at the following numbers:

- | | |
|---|----------------|
| • Statistical Information Service | 1-800-263-1136 |
| • National telecommunications device for the hearing impaired | 1-800-363-7629 |
| • Fax line | 1-514-283-9350 |

Standards of service to the public

Statistics Canada is committed to serving its clients in a prompt, reliable and courteous manner. To this end, Statistics Canada has developed standards of service that its employees observe. To obtain a copy of these service standards, please contact Statistics Canada toll-free at 1-800-263-1136. The service standards are also published on www.statcan.gc.ca under “Contact us” > “[Standards of service to the public](#).”

Note of appreciation

Canada owes the success of its statistical system to a long-standing partnership between Statistics Canada, the citizens of Canada, its businesses, governments and other institutions. Accurate and timely statistical information could not be produced without their continued co-operation and goodwill.

Published by authority of the Minister responsible for Statistics Canada

© His Majesty the King in Right of Canada, as represented by the Minister of Industry, 2024

All rights reserved. Use of this publication is governed by the Statistics Canada [Open Licence Agreement](#).

An [HTML version](#) is also available.

Cette publication est aussi disponible en français.

Design-based conformal prediction

Jerzy Wiecek¹

Abstract

Conformal prediction is an assumption-lean approach to generating distribution-free prediction intervals or sets, for nearly arbitrary predictive models, with guaranteed finite-sample coverage. Conformal methods are an active research topic in statistics and machine learning, but only recently have they been extended to non-exchangeable data. In this paper, we invite survey methodologists to begin using and contributing to conformal methods. We introduce how conformal prediction can be applied to data from several common complex sample survey designs, under a framework of design-based inference for a finite population, and we point out gaps where survey methodologists could fruitfully apply their expertise. Our simulations empirically bear out the theoretical guarantees of finite-sample coverage, and our real-data example demonstrates how conformal prediction can be applied to complex sample survey data in practice.

Key Words: Conformal prediction; Machine learning; Cross validation; Predictive modeling; Complex sample survey designs.

1. Introduction

What is conformal prediction? Conformal prediction, also called conformal inference, is a family of general-purpose approaches to constructing prediction intervals or prediction sets, which can be wrapped around almost any predictive modeling algorithm. Specifically, imagine that we have fit a function \hat{f}_n to a set of training data (X_i, Y_i) for $i = 1, \dots, n$ that allows us to make a point prediction for a new observation's response value Y_{n+1} when \hat{f}_n is evaluated at the covariate value X_{n+1} , and our goal is to report a level $1 - \alpha$ prediction set for Y_{n+1} . Let us require only that all $n + 1$ of the points (X_i, Y_i) are an exchangeable sample from some common distribution P , and that the algorithm used to fit \hat{f}_n treats these points symmetrically (which rules out e.g. algorithms that give more weight to more-recent data over time). (A sequence of random variables is exchangeable when its joint probability distribution would not change if you permuted the random variables. More precisely, a sequence X_1, X_2, \dots is exchangeable if for each n and each permutation π of $\{1, \dots, n\}$, the joint distribution of $(X_{\pi(1)}, \dots, X_{\pi(n)})$ is just the same as the distribution of (X_1, \dots, X_n) (Durrett, 2019). Examples of exchangeability include iid sequences as well as simple random sampling (with or without replacement).) If so, we can use standard conformal prediction methods to construct \hat{C}_n which are either prediction intervals (for a regression problem) or prediction sets (for a classification problem), such that $\mathbb{P}[Y_{n+1} \in \hat{C}_n(X_{n+1})] \geq 1 - \alpha$, where the probability is taken over repeated sampling of all $n + 1$ points. In Section 2 we describe two standard approaches to conformal prediction and describe the intuition behind how they work.

Although survey data analysis has traditionally focused on estimation and testing, predictive modeling with complex sample survey data is not uncommon. To name just a few examples: Hong and He (2010) use the Second Longitudinal Study on Aging (National Center for Health Statistics, 2016) to fit a predictive

1. Jerzy Wiecek, Department of Statistics, Colby College, Waterville, Maine, USA. E-mail: jawieczo@colby.edu.

model that can be used to monitor functional mobility status among the elderly. Kshirsagar, Wieczorek, Ramanathan and Wells (2017) use 2015 Living Conditions Monitoring Survey data (Central Statistical Office, Zambia, 2015) to build a model to classify households by poverty level, at a range of poverty thresholds. Krebs, Reeves and Baggett (2019) use Forest Inventory and Analysis data (Bechtold and Patterson, 2005) to build random-forest models that can predict understory vegetation structure, intending for these models to be implemented in the Forest Vegetation Simulator (Crookston and Dixon, 2005). In each case, there is interest in using the model to make unit-level predictions, not just to estimate regression coefficients or population means. Yet, with the exception of Hong and He (2010), many such papers have not been able to provide prediction intervals or sets.

Why is this interesting to survey methodologists? For practitioners, survey data analysts who build predictive models – like those mentioned above – could apply conformal methods to provide design-based prediction intervals or sets (Section 3). Such prediction intervals or sets would have guaranteed coverage, even for novel machine learning algorithms. Conformal methods may not add much value to older methods like linear regression, which already has well-known Gaussian-based prediction intervals and has already been adapted to account for complex survey designs. But for newer methods – such as prediction algorithms that have not yet been adapted to work optimally with survey designs, or non-probabilistic algorithms that do not come with “built-in” prediction intervals such as the simulation models of LeRoy and Schafer (2021) – we can apply conformal prediction methods that guarantee coverage based on the sampling design alone. Those who collect and pre-process the survey data may also find uses for conformal methods, for instance in nonresponse prediction, imputation, or data cleaning (Section 5).

More philosophically, the principles behind conformal methods align with complex survey sampling inference. [a] Conformal methods and many design-based methods provide exact, finite-sample coverage guarantees, not asymptotic or approximate guarantees. [b] Unlike traditional model-based approaches (such as Gaussian-errors prediction intervals for regression), conformal prediction does not require assumptions about the distribution P – only exchangeable sampling. Likewise, design-based methods do not require assumptions about how the data values are distributed in the population – only knowledge of the sampling design. Finally, [c] conformal guarantees hold even if the predictive model is not a good fit to P (though of course the prediction intervals or sets may be larger then), and many model-assisted methods in design-based inference also have guarantees that hold whether or not the model is a good fit to the population.

Hence, we believe this is a good opportunity to cross-pollinate ideas between the survey methodology and machine learning research communities. The growth of conformal prediction indicates that machine learning practitioners are seeking methods that are guaranteed to work based only on the sampling design, not on the distribution of the data values. Survey methodologists, having the right expertise to meet that demand, may enjoy applying their skillset to this new challenge.

Why is this not already in use for survey sampling? Despite the apparent affinities, conformal methods have not previously been studied under a complex sample survey framework. Conformal prediction has

been under development for several decades, initially by Vladimir Vovk and colleagues (Vovk, Gammerman and Shafer, 2022) and more recently by a wide range of statistical and machine learning researchers. Conformal methods are starting to be deployed in real-world settings, such as the Washington Post’s 2020 presidential election tracker, which reported prediction intervals of the votes for each party that were updated in real time as voting districts slowly reported their results (Cherian and Bronner, 2020). Yet the requirement of exchangeability has made traditional conformal methods inapplicable for complex sampling designs other than simple random sampling (with or without replacement).

However, recently Tibshirani, Barber, Candes and Ramdas (2019) extended conformal prediction to what they call “weighted exchangeable” sampling, while Dunn, Wasserman and Ramdas (2022) derived several “hierarchical” conformal methods that can apply to cluster sampling. The present work builds on these two papers to begin providing conformal guarantees for complex sampling designs.

Briefly, instead of assuming that all $n+1$ data points are exchangeable, we will assume that the n training points came from a known complex sampling design on a finite population, and the test point was selected uniformly at random from the same finite population. The latter assumption is simply a way of phrasing the guarantee of marginal coverage across the finite population. In Section 3, we summarize the complex survey settings in which such conformal prediction methods are currently known to have coverage guarantees.

1.1 Related work

Prediction intervals or tolerance regions have a long history in statistics; for a recent review, see Tian, Nordman and Meeker (2022). Section 13.3.3 of Vovk et al. (2022) connects conformal methods to much of this earlier work. However, conformal prediction is currently a broad and very active research area, and we encourage interested readers to look into the following general resources: Angelopoulos and Bates (2022) provide an introductory review article; Vovk and colleagues maintain a list of their own recent work at <http://alrw.net/>; Manokhin (2022) maintains a frequently updated list of conformal papers, tutorials, and software; and COPA, the Symposium on Conformal and Probabilistic Prediction with Applications, is an annual conference on conformal prediction and its extensions: <https://copa-conference.com/>.

For dependent data, our paper builds largely on the work of Dunn et al. (2022) and Tibshirani et al. (2019). Fong and Holmes (2021) give a distinct, Bayesian approach to conformal prediction for hierarchical data. Barber, Candès, Ramdas and Tibshirani (2022) relax the usual conformal-prediction requirement that the fitted prediction model must treat data points symmetrically, and Fannjiang, Bates, Angelopoulos, Listgarten and Jordan (2022) address “feedback covariate shift”, where the choice of which test data to sample depends on results from the training data, which could be useful in addressing adaptive sampling for surveys (Thompson, 1997). Conformal methods for nonstationarity over time (Chernozhukov, Wüthrich and Yinchu, 2018; Oliveira, Orenstein, Ramos and Romano, 2022), including distribution drift (Gibbs and

Candes, 2021), would be useful in developing conformal methods for panel surveys. Recently, Lunde (2023) developed conformal methods for network data under non-uniform sampling, with implications for respondent-driven sampling in surveys.

Few papers have applied conformal methods to survey samples. Bersson and Hoff (2022) do use conformal prediction for small area estimation problems which rely on survey data. However, they work in a model-based framework and assume that sampled units are exchangeable within each small area, rather than focusing on design-based inference under a complex sampling design. Romano, Patterson and Candes (2019) and several followup papers (Sesio and Candes, 2020; Sesio and Romano, 2021; Feldman, Bates and Romano, 2021; Bai, Mei, Wang, Zhou and Xiong, 2022) illustrate conformal prediction methods using unit-level records from a complex survey: the Medical Expenditure Panel Survey (Agency for Healthcare Research and Quality, 2017). However, they appear to ignore the documented sampling design and treat it as exchangeable.

Apart from conformal methods, there has been increasing interest in applying machine learning methods to complex survey data. Dagdoug, Goga and Haziza (2022) summarize the current state of the art for using survey data with many machine learning algorithms: penalized regression (McConville, Breidt, Lee and Moisen, 2017), k nearest neighbors (Yang and Kim, 2019), random forests (Dagdoug, Goga and Haziza, 2021), and others. However, along with Sande and Zhang (2021), they focus on estimation of means and totals, rather than on predicting individual response values. Even when such algorithms have been adapted to make individual predictions that account for the survey design, many still do not offer procedures for creating prediction intervals, which design-based conformal methods could provide.

1.2 Our contributions

We introduce survey methodologists to key definitions and intuitions behind standard conformal methods under exchangeable sampling as well as their extensions to complex sampling (Section 2). Furthermore, we derive exact, finite-sample, design-based coverage guarantees for applying conformal inference to data from several fundamental sampling designs, building on results from the weighted exchangeable or hierarchical frameworks (Section 3).

Next, we illustrate our design-based guarantees through simulations and a real-data example (Section 4). We show that results can differ on real data depending on whether or not conformal methods account for the survey design, and we show how conformal methods are affected by different sampling designs. We then discuss several practical considerations and future challenges to be addressed in applying conformal methods to survey data (Section 5). We also suggest ways that conformal methods might be useful in conducting surveys. Finally, we invite survey researchers to contribute to the literature on conformal methods (Section 6). We anticipate that advances from the design-based perspective will turn out to be useful to the general community of conformal researchers.

2. Introduction to conformal inference

2.1 Definitions

Because one of our main contributions is a corollary of results in Tibshirani et al. (2019), we restate without proof some of their notation, definitions, and results here and in Section 3.

Let $\text{Quantile}(\beta; F)$ denote the level β quantile of distribution F , so that for $Y \sim F$,

$$\text{Quantile}(\beta; F) = \inf \{y : \mathbb{P}\{Y \leq y\} \geq \beta\}.$$

We allow for distributions on the augmented real line, $\mathbb{R} \cup \{\infty\}$. We use $v_{1:n} = \{v_1, \dots, v_n\}$ to denote a multiset, meaning that it is unordered and can allow the same element to appear several times. We use δ_a to denote a point mass at the value a . If $v_{1:n}$ is an exchangeable sample, its empirical probability measure is $n^{-1} \sum_{i=1}^n \delta_{v_i}$, and the level β quantile of its empirical distribution is $\text{Quantile}(\beta; v_{1:n})$ which is the $\lceil \beta n \rceil$ smallest value in $v_{1:n}$, where $\lceil \cdot \rceil$ is the ceiling function.

Now we can state the general-purpose “quantile lemma” that forms the foundation for conformal inference in the exchangeable setting. For instance, if we apply this lemma directly to data from an exchangeable sample of scalar random variables V_1, \dots, V_n , we obtain a level β one-sided prediction interval for a new observation V_{n+1} .

Lemma 1 (Tibshirani et al. (2019)). *If V_1, \dots, V_{n+1} are exchangeable random variables, then for any $\beta \in (0, 1)$, we have*

$$\mathbb{P}\{V_{n+1} \leq \text{Quantile}(\beta; V_{1:n} \cup \infty)\} \geq \beta.$$

Furthermore, if ties between V_1, \dots, V_{n+1} occur with probability zero, then the above probability is upper bounded by $\beta + 1/(n+1)$.

In order to use this lemma to do conformal prediction for nontrivial regression or classification problems, we must also choose a *score function* \mathcal{S} , which takes the following arguments: a point (x, y) , and a multiset Z (meaning that \mathcal{S} must treat the points in Z as unordered). The points in Z will typically be the n observations $Z_i = (X_i, Y_i)$ for $i = 1, \dots, n$, possibly along with one additional hypothetical observation or test case; (x, y) is typically one of the points in Z . \mathcal{S} should return a real value, such that lower values indicate that (x, y) “conforms” to Z better. Finally, use \mathcal{S} to define *nonconformity scores*

$$V_i^{(x,y)} = \mathcal{S}(Z_i, Z_{1:n} \cup \{(x, y)\}), i = 1, \dots, n, \text{ and } V_{n+1}^{(x,y)} = \mathcal{S}((x, y), Z_{1:n} \cup \{(x, y)\}). \quad (2.1)$$

The reason we require Z to be an unordered multiset is to ensure that if the observations (X_i, Y_i) are exchangeable, then the scores $V_i^{(x,y)}$ will be too.

For instance, in a regression context, we might choose $\mathcal{S}((x, y), Z) = |y - \hat{f}(x)|$ where \hat{f} is some regression function fitted using all of Z , including (x, y) . If the absolute-residual nonconformity score

$V_{n+1}^{(x,y)}$ is small relative to all of the scores $V_i^{(x,y)}$, this suggests that (x, y) conforms well to the overall trend in Z .

In a classification setting, we might choose $\mathcal{S}((x, y), Z) = 1 - \hat{f}(x)_y$ where \hat{f} is the probability of class y estimated by some classification function fitted using all of Z , again including (x, y) . A small value of this nonconformity score $V_{n+1}^{(x,y)}$ suggests that (x, y) conforms well to Z , in the following sense: This observation with covariates x has a high estimated probability of being in class y based on trends in Z , and it does indeed belong to class y .

We will also use following common abbreviations: SRS (simple random sampling), WR (with replacement), WOR (without replacement), PPS (probability proportional to size sampling). N denotes a population size and n typically denotes a sample size, except for the “split conformal” methods (defined in Section 2.4) which split the data into a “proper training set” of size m and a “calibration set” of size n .

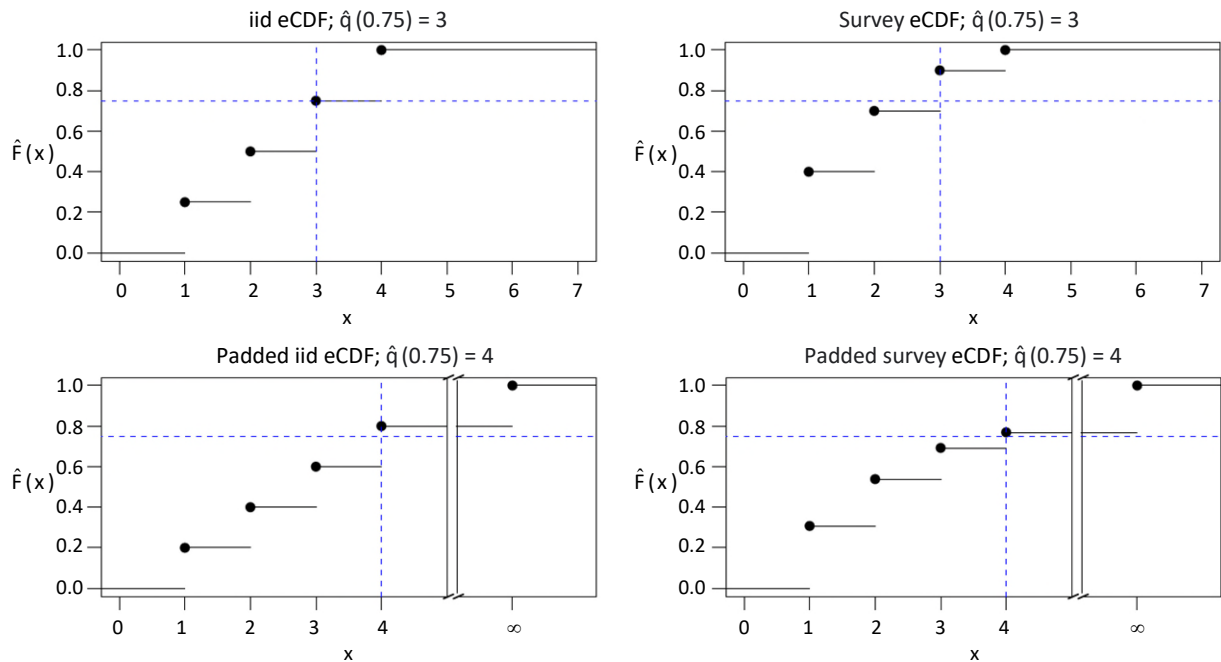
2.2 Intuition under exchangeable sampling

Quantile lemma: The lower bound in Lemma 1 has a simple rationale. First, let \hat{q}_{n+1} be the $\lceil \beta(n+1) \rceil$ smallest value in $V_{1:(n+1)}$, and let \hat{q}_{conf} be the $\lceil \beta(n+1) \rceil$ smallest value in $V_{1:n} \cup \infty$. By exchangeability, the rank of V_{n+1} among all the V_i is uniformly distributed over $\{1, \dots, n+1\}$, so V_{n+1} is at or below \hat{q}_{n+1} with probability exactly β . Then since $\hat{q}_{\text{conf}} \geq \hat{q}_{n+1}$, V_{n+1} is at or below \hat{q}_{conf} with probability *at least* β . Note that these probabilities are marginal (over exchangeable sampling of all the V_1, \dots, V_{n+1} together) – not conditional on the first n observations nor conditional on the test observation.

This is an exact finite-sample result that only relies on exchangeability. By contrast, if we used instead \hat{q}_n , the $\lceil \beta n \rceil$ smallest value in $V_{1:n}$, then it may be an asymptotically good estimate of the population quantile, but it would require stronger conditions on the data distribution. Even then, the probability that V_{n+1} is at or below \hat{q}_n would be only *approximately* β .

To see why $\hat{q}_{\text{conf}} \geq \hat{q}_{n+1}$, imagine an empirical cumulative distribution function (eCDF) plot of just the first n datapoints, with step heights of $1/n$ at each observed value. (See the top left subplot of Figure 2.1.) If we add one more datapoint, the step heights of the eCDF will be $1/(n+1)$ instead, and the lower β quantile will be \hat{q}_{n+1} as above. It may be larger or smaller than the lower β quantile of the first n datapoints, depending on how large the last datapoint is. The lower β quantile is largest if we choose a $(n+1)^{\text{th}}$ value that is larger than any of the first n observed values, say at ∞ . That is because if the added datapoint is larger than any of the others, it pushes the rest of the eCDF down, so the horizontal line with y -intercept of β crosses the eCDF at an x value further to the right. In this most extreme case, the lower β quantile is \hat{q}_{conf} as above, and therefore $\hat{q}_{\text{conf}} \geq \hat{q}_{n+1}$. (See the bottom left subplot of Figure 2.1.) By contrast, if instead of ∞ we had chosen to add a value smaller than some of the first n observations, it would have pushed part of the eCDF upward and may have caused the lower β quantile to become smaller than \hat{q}_{n+1} .

Figure 2.1



Notes: Top left: eCDF for an iid sample with values $\{1, 2, 3, 4\}$. At each value there is a vertical jump of $1/n = 0.25$. The 75th percentile is at 3.
 Bottom left: same eCDF but padded with an extra value at ∞ . Now at each value there is a vertical jump of $1/n + 1 = 0.20$. The 75th percentile is between 3 and 4, so we round up to 4.
 Top right: eCDF for a survey sample with values $\{1, 2, 3, 4\}$ and corresponding survey weights $\{4, 3, 2, 1\}$. At each value there is a vertical jump proportional to its survey weight. The 75th percentile is between 2 and 3, so we round up to 3.
 Bottom right: same eCDF but padded with an extra value at ∞ , corresponding to a not-yet-observed unit in the population that would have a survey weight of 3 if it were to be sampled. The vertical jumps at each value are still proportional to the survey weights, but now rescaled to make room for the weight of the extra value at ∞ . The 75th percentile is between 3 and 4, so we round up to 4.

Put another way, the probability that $\{\text{we take an exchangeable sample of size } n + 1 \text{ in which the last observation is greater than the } \lceil (1 - \alpha)(n + 1) \rceil \text{ smallest of the first } n \text{ observations}\}$ is at most α . This is a strict inequality, not approximate or asymptotic.

Regression prediction intervals: Next, we can apply this quantile lemma in a regression setting. (In Section 2.5 we describe an approach for classification problems.) Imagine that we are given a fixed regression function \hat{f} for making predictions of some real-valued random variable Y from some covariate vectors X for data from this population. (For the moment, assume \hat{f} is not data-dependent; perhaps it was chosen *a priori*.) Also imagine that we have some residuals from applying \hat{f} to a “calibration set” of n different data points $(X_1, Y_1), \dots, (X_n, Y_n)$. (We recognize that the term “calibration” traditionally has a completely different meaning in survey sampling. In this paper, we will use “calibration” only in this conformal inference sense.) Finally, we have one other covariate vector X_{n+1} , drawn exchangeably from the same population, but we do not know its Y value. We want to get a level $(1 - \alpha)$ prediction interval for Y_{n+1} at this test-case X_{n+1} .

We can simply apply the “quantile lemma” logic to the n absolute residuals from the calibration set. These absolute residuals are our nonconformity scores. If we let \hat{q} be the $\lceil (1-\alpha)(n+1) \rceil$ smallest absolute residual, then $\hat{C}(X_{n+1}) = \hat{f}(X_{n+1}) \pm \hat{q}$ is a prediction interval for Y_{n+1} with guaranteed coverage of at least $1-\alpha$.

Marginal vs. conditional coverage: This coverage is marginal across all samples of size $(n+1)$, with a size- n calibration set plus one new test case. It is not conditional on the specific test case X_{n+1} we chose; it assumes that both the calibration set and the test case together are exchangeable. Nonetheless, it is an exact finite-sample result. Barber, Candès, Ramdas and Tibshirani (2021) show that although no conformal method could generally guarantee “exact conditional coverage” (conditioning on the exact value of X), certain relaxed versions of conditional coverage are achievable. Angelopoulos and Bates (2022) review approaches to assessing and controlling various forms of conditional coverage, noting that marginal coverage alone may be insufficient, e.g. if it happens to be achieved by low coverage in a rare but important subpopulation and high coverage elsewhere. In Appendix A we review links between conditional coverage and the related idea of “adaptive” prediction regions.

2.3 Intuition under complex sampling

When the data are not exchangeable, Tibshirani et al. (2019) have extended conformal prediction to a setting called “covariate shift”, in which the distribution of X is different for the training and/or calibration sets than for the test set, but the conditional distribution of $Y|X$ remains the same. They define a condition called “weighted exchangeability” and show how to make conformal inference work for such data. In the present paper, we will show that certain classic finite-population sampling designs can be treated as a special case of covariate shift, and therefore the results of Tibshirani et al. (2019) apply.

Specifically, imagine we sample n cases with replacement from a finite population with known but unequal sampling probabilities; for instance, we may be using PPS sampling. Assume also that we take a SRS of just one test case from the full finite population. Now, we need a “survey weighted quantile lemma” that tells us how to find an adjusted quantile \hat{q} of the complex sample, such that a test case nonconformity score is no larger than \hat{q} with probability at least $1-\alpha$. We will plug the n complex sample observations from our calibration set into a regression function \hat{f} and find \hat{q} for the absolute residuals. Then the probability of our test case (X_{n+1}, Y_{n+1}) having a larger absolute residual is at most α , and so $\hat{f}(X_{n+1}) \pm \hat{q}$ is a level $1-\alpha$ prediction interval for Y_{n+1} .

In Section 3 we will show that in this unequal-probabilities setting, it is enough to mimic the exchangeable setting, except that instead of using an equal-weights eCDF to get the quantiles, we use a survey-weighted eCDF; see e.g. Section 5.11 of Särndal, Swensson and Wretman (1992). The additional observation at ∞ will simply be assigned the known sampling weight that the test case (X_{n+1}, Y_{n+1}) would have had, under the complex sampling design used to sample the first n units. See the right half of Figure 2.1 for an illustration.

We acknowledge that there is a long-running debate in the survey sampling literature about whether and how to use sampling weights for model fitting and inference (Fienberg, 2010; Lumley and Scott, 2017). Our goal in the present paper is not to take a stance in this debate, but simply to show how conformal inference can be applied *if* the data analyst is taking a design-based perspective. In a model-based analysis, if the design features can justifiably be ignored, standard conformal inference methods may be used.

2.4 Split vs. full conformal

Above, for simplicity we have assumed that a \hat{f} has been provided for us. More typically, we will need to fit \hat{f} using data from the sample at hand. In the “split conformal” approach (Lei, G’Sell, Rinaldo, Tibshirani and Wasserman, 2018), also called “inductive conformal” (Papadopoulos, Proedrou, Vovk and Gammernan, 2002), we start with an exchangeable dataset of $m+n$ observations, and we split them at random into a “proper training set” of size m to be used for training \hat{f} , plus a calibration set of size n as described above. In this situation, coverage is still marginal over the calibration set plus one test case, but now it is conditional on the proper training set.

An optimal sample splitting ratio m/n for split conformal is not known. Larger m/n should lead to better estimates of \hat{f} , and therefore shorter conformal prediction interval lengths on average. On the other hand, larger m/n also leads to a smaller calibration set, and therefore more-variable conformal prediction interval lengths.

If we do not wish to lose statistical efficiency by splitting our data, we can use a more computationally-intensive “full conformal” approach (Vovk et al., 2022). In that approach, we no longer need a separate calibration set, and we let n denote the total number of our complete-data training cases. We also have one test case with only the covariates X_{n+1} known. Then we repeat the following process for many y values:

- Choose a hypothetical response value $y \in \mathbb{R}$.
- Fit \hat{f}_y to an augmented dataset in which we pretend this y -value is correct: $(X_1, Y_1), \dots, (X_n, Y_n), (X_{n+1}, y)$. Find the $n+1$ nonconformity scores, e.g. the absolute residuals: $R_{y,i} = |Y_i - \hat{f}_y(X_i)|$ for $i=1, \dots, n$ and $R_{y,n+1} = |y - \hat{f}_y(X_{n+1})|$. Find their $1-\alpha$ quantile: \hat{q}_y is the $\lceil (1-\alpha)(n+1) \rceil$ smallest value of $R_{y,i}$ for $i=1, \dots, n+1$.
- If $R_{y,n+1} \leq \hat{q}_y$, we say y “conforms” to the rest of the data, and we add y to our prediction interval.

In terms of notation, refer back to (2.1). Full conformal nonconformity scores are calculated using $Z = Z_{1:n} \cup \{(X_{n+1}, y)\}$ and refitting a new prediction model \hat{f}_y for each new y or X_{n+1} . By contrast, split conformal nonconformity scores use a fixed \hat{f} conditional on the proper training set. Each of the first n nonconformity scores is calculated only using that data point (X_i, Y_i) and the fixed \hat{f} , and there is no need to calculate a $(n+1)^{\text{th}}$ nonconformity score. Under the notation in (2.1), we allow Z to be ignored and $V_{n+1}^{(x,y)}$ to be undefined for split conformal.

Compared to split conformal, the full conformal procedure reduces the variance in the reported prediction interval endpoints. On the other hand, full conformal is far more computationally intensive than split conformal, as \hat{f}_y has to be refit for every candidate y at a given test point X_{n+1} , and the entire interval needs to be refit for every test point.

In empirical comparisons, Lei et al. (2018) found that split conformal and full conformal often produce very similar prediction intervals. They recommend using split conformal, on the grounds that it is faster to compute with little loss of efficiency, although full conformal avoids randomness in the data split.

2.5 Classification problems and prediction sets

For brevity, we focus on regression prediction intervals, but classification prediction sets are another common use case. For instance, consider split conformal inference for a multi-class probabilistic classifier. Examples range from simple logistic regression to deep neural networks. First, fit the classifier to the proper training set to get a function $\hat{f}(x)_y$, whose outputs are the estimated probability of class y for an input x . Then, using the calibration set, each calibration “residual” or nonconformity score V_i is calculated as $1 - \hat{f}(X_i)_y$ for each calibration observation $i = 1, \dots, n$. Next, find the corrected $1 - \alpha$ quantile of these probabilities: let \hat{q} be the $\lceil (1 - \alpha)(n + 1) \rceil$ smallest value of V_i for $i = 1, \dots, n$. Finally, for a new test case X_{n+1} , find its estimated probability for each class y , and choose all the classes whose nonconformity scores are below the corrected quantile: $\hat{C}(X_{n+1}) = \{y : 1 - \hat{f}(X_{n+1})_y \leq \hat{q}\}$. Romano, Sesia and Candes (2020) and Angelopoulos and Bates (2022) discuss refinements to this approach.

3. Methods

For the reader’s convenience, we restate several of the key results from Tibshirani et al. (2019) below. Next we prove that we can apply these results to unequal-probability sampling with replacement. We follow by discussing methods for sampling without replacement, cluster sampling based on Dunn et al. (2022), stratified sampling, and post-stratification.

Most of the methods below apply both to full and split conformal. However, for split conformal, we recommend using the design-based approach of Wieczorek, Guerin and McMahon (2022) to split survey data into a proper training set and a calibration set. This way, both the proper training set and the calibration set will mimic the original sampling design, which ensures that the methods below are safe to apply to your calibration set. By contrast, a simple random split can cause the calibration set to have different properties than a clustered or stratified sampling design.

3.1 Previous results for covariate shift (Tibshirani et al., 2019)

Tibshirani et al. (2019) define the “covariate shift” setting as follows:

$$(X_i, Y_i) \stackrel{\text{iid}}{\sim} P = P_X \times P_{Y|X}, i = 1, \dots, n,$$

$$(X_{n+1}, Y_{n+1}) \sim \tilde{P} = \tilde{P}_X \times P_{Y|X}, \text{ independently.} \tag{3.1}$$

Note that the conditional distribution of $Y|X$ remains the same as the marginal distribution of X changes.

Assuming that P_X and \tilde{P}_X are known, we can define likelihood ratio weight functions $w = d\tilde{P}_X/dP_X$. We use these to define a second set of weights:

$$p_i^w(x) = \frac{w(X_i)}{\sum_{j=1}^n w(X_j) + w(x)}, \quad i = 1, \dots, n, \quad \text{and} \quad p_{n+1}^w(x) = \frac{w(x)}{\sum_{j=1}^n w(X_j) + w(x)}. \tag{3.2}$$

In this setting, we can state a weighted, nonexchangeable counterpart to Lemma 1. Although Tibshirani et al. (2019) give and prove a more general version, here we only state a version tailored to the covariate shift setting.

Lemma 2. *Assume data from the model (3.1). Assume \tilde{P}_X is absolutely continuous with respect to P_X , and denote $w = d\tilde{P}_X/dP_X$. For any $\beta \in (0, 1)$,*

$$\mathbb{P} \left\{ V_{n+1} \leq \text{Quantile} \left(\beta; \sum_{i=1}^n p_i^w(x) \delta_{V_i} + p_{n+1}^w(x) \delta_\infty \right) \right\} \geq \beta,$$

where $V_i^{(x,y)}$, $i = 1, \dots, n + 1$ are as defined in (2.1), and p_i^w , $i = 1, \dots, n + 1$ are as defined in (3.2).

Proof. Apply Lemma 3 of Tibshirani et al. (2019) in the covariate shift setting of (3.1).

This weighted quantile lemma allows conformal inference in the covariate shift setting.

Corollary 1 (Tibshirani et al. (2019)). *Assume data from the model (3.1). Assume \tilde{P}_X is absolutely continuous with respect to P_X , and denote $w = d\tilde{P}_X/dP_X$. For any score function \mathcal{S} , and any $\alpha \in (0, 1)$, define for $x \in \mathbb{R}^d$,*

$$\hat{C}_n(x) = \left\{ y \in \mathbb{R} : V_{n+1}^{(x,y)} \leq \text{Quantile} \left(1 - \alpha; \sum_{i=1}^n p_i^w(x) \delta_{V_i^{(x,y)}} + p_{n+1}^w(x) \delta_\infty \right) \right\},$$

where $V_i^{(x,y)}$, $i = 1, \dots, n + 1$ are as defined in (2.1), and p_i^w , $i = 1, \dots, n + 1$ are as defined in (3.2). Then \hat{C}_n satisfies $\mathbb{P} \{ Y_{n+1} \in \hat{C}_n(X_{n+1}) \} \geq 1 - \alpha$.

This is the weighted “full conformal” approach. For the “split conformal” approach, we restate part of Section A.3 from the supplement to Tibshirani et al. (2019). Let $(X_1^0, Y_1^0), \dots, (X_m^0, Y_m^0)$ be a proper training set of size m , used for fitting the regression function \hat{f}_0 . Also let $(X_1, Y_1), \dots, (X_n, Y_n)$ be the calibration set of size n and let (X_{n+1}, Y_{n+1}) be the test case. Then weighted split conformal prediction is a special case of Corollary 1 in which \hat{f}_0 is treated as fixed, and e.g. if we use the absolute-residual score function, the prediction interval simplifies to

$$\hat{C}_n(x) = \hat{f}_0(x) \pm \text{Quantile}\left(1-\alpha; \sum_{i=1}^n p_i^w(x) \delta_{|Y_i - \hat{f}_0(x_i)|} + p_{n+1}^w(x) \delta_\infty\right),$$

with weights p_i^w as defined in (3.2). By Corollary 1, this has coverage at least $1-\alpha$ conditional on the proper training set $(X_1^0, Y_1^0), \dots, (X_m^0, Y_m^0)$. Similar results apply for other score functions and for classification problems.

Finally, Remark 3 of Tibshirani et al. (2019) notes that the results above still hold if the likelihood ratio weights have an unknown normalization constant, i.e. if $w \propto d\tilde{P}_X/dP_X$, because this constant cancels out in the final weights in (3.2).

3.2 Unequal probability sampling with replacement

Now, consider independently sampling WR from a finite population where each unit has unequal but known sampling probabilities, such as in PPS sampling. Assume no stratification, clustering, or other design constraints. Our goal is to use this survey sampling design to get prediction intervals that cover Y for most of the N units in the finite population.

Below, we show that this is a special case of the covariate shift setting of Tibshirani et al. (2019). Their distribution P_X is replaced with the sampling probabilities for each unit in the finite population, and \tilde{P}_X is replaced with a uniform distribution over the same population units. The only randomness is in the sampling; all covariates and response variables are fixed in the population, as is usual in design-based inference.

Let j index the universe of population units from 1 to N . Once a sample S is selected, it can be written as a length- n vector of sampled unit IDs, so that the elements of S take values in $\{1, \dots, N\}$. Often, survey researchers find it easier to work with the random vector S by rewriting it as a length- N vector of sampling indicators $Z_j = \{1 \text{ if } j \in S, \text{ and } 0 \text{ otherwise}\}$. However, in the covariate shift setting, it will be easier to work with S directly. While using j to index the population units, we use i for indexing the elements of S . For example, $S_1 = 4$ would mean that the first unit we sampled was the fourth member of the population.

Lemma 3. *Assume our training sample is a sample WR of size n where each unit is drawn independently from a finite population of size N , with possibly-unequal but nonzero and known sampling probabilities π_j for each population ID $j \in 1, \dots, N$. Also assume our test case is a single observation sampled uniformly at random from the population. Let S be a vector whose first n elements are the population IDs of our n training cases, and whose last element is the population ID for our test case.*

- *Let our finite population consist of N units, each with its own fixed data vector (X_j, Y_j) and sampling probability π_j for $j = 1, \dots, N$. The conditional distribution $P_{(X,Y)|S}$ is a deterministic lookup table: Once we have sampled a random S_i for some $i \in \{1, \dots, n+1\}$, we observe its corresponding covariate and response values (X_{S_i}, Y_{S_i}) .*
- *Let the training distribution P_S correspond to our complex survey design, which consists of n iid draws from a categorical distribution whose categories are simply the unit IDs $j \in \{1, \dots, N\}$, with known probabilities π_j . (The sampling indicator vector Z is a draw from a multinomial*

distribution (sampling WR) or a multivariate hypergeometric distribution (sampling WOR), where the j^{th} element counts how often that unit was selected. Sampling from a categorical distribution is identical, but instead of recording counts, we record each of the selected unit IDs individually (Tu, 2014.)

- Let the “covariate-shifted” test distribution \tilde{P}_S consist of one draw from the same set of unit IDs as the training distribution, but with uniform probability: let $\tilde{\pi}_j = 1/N$ for all $j = 1, \dots, N$.

Then this is a special case of the covariate shift setting, with S and (X, Y) playing the roles that X and Y respectively played in (3.1):

$$\begin{aligned} (S_i, (X_{S_i}, Y_{S_i})) &\stackrel{\text{iid}}{\sim} P = P_S \times P_{(X, Y)|S}, i = 1, \dots, n, \\ (S_{n+1}, (X_{S_{n+1}}, Y_{S_{n+1}})) &\sim \tilde{P} = \tilde{P}_S \times P_{(X, Y)|S}, \text{ independently,} \end{aligned} \quad (3.3)$$

and $w(S_i) \equiv 1/\pi_{S_i}$.

Proof. In the training data sampling design, units are sampled with replacement from a fixed population, and thus they are iid. Even though individual population units have different sampling probabilities, each unit in a training sample is drawn from the same categorical distribution. Also, because each (X_j, Y_j) is paired with a fixed j for all $j = 1, \dots, N$, the conditional distribution $P_{(X, Y)|S}$ is the same for training and test data. Hence, these training and test distributions match the requirements in (3.1).

Furthermore, since π_j is nonzero for all population units, we have $w = d\tilde{P}_S/dP_S = (1/N)/\pi_{S_i} \propto 1/\pi_{S_i}$. By Remark 3 of Tibshirani et al. (2019), it is safe to ignore the normalization constant and set $w(S_i) = 1/\pi_{S_i}$ directly.

Corollary 2. Assume that n training cases and one test case are drawn as described in Lemma 3. Then Lemma 2, the weighted full conformal results from Corollary 1, and the weighted split conformal results from Section A.3 of Tibshirani et al. (2019) hold, with p_i^w defined by using the inverse-probability sampling weights $1/\pi_{S_i}$ for the likelihood ratio weight function w in (3.2).

Proof. In this setting, P_S and \tilde{P}_S are discrete distributions with identical support, so \tilde{P}_S is absolutely continuous with respect to P_S . All other conditions of Corollary 1 and Section A.3 of Tibshirani et al. (2019) are met by Lemma 3.

As suggested by the intuition in Section 2.3, we carry out conformal inference by replacing the exchangeable eCDF with a survey-weighted eCDF in which the first n observations (the training sample) have their usual inverse-probability sampling weights. The $(n+1)^{\text{th}}$ sample (the test case) is assigned the known sampling weight $1/\pi_{S_{n+1}}$ that population unit S_{n+1} would have had under the original sampling design. The assumption of uniform sampling of test cases simply lets us guarantee marginal coverage across the entire finite population. We do need to know what π_j would have been for every unit in the population,

which may be reasonable for the organization carrying out the survey but not for end users of the data; Section 5.1 discusses ways to address this issue.

In a situation with no covariates, we can replace X with a constant in Lemma 3, and apply Lemma 2 to get a prediction interval for Y that is not conditional on X , although it may depend on $\pi_{S_{n+1}}$.

3.3 Unequal probability sampling without replacement

SRSWOR is exchangeable, so the usual conformal methods apply directly. But other kinds of sampling WOR are not a special case of the covariate shift setting above, because the training data are no longer independent. While Tibshirani et al. (2019) do provide more general results under a relaxed condition they call “weighted exchangeability”, it is not immediately clear that this condition can account for sampling WOR.

However, if $n \ll N$, then statistical properties derived under sampling WR are often fairly good approximations to the actual properties under assuming sampling WOR. Our simulations in Section 4.2, in which we sample WOR, suggest that this is likely to hold true for conformal methods as well.

3.4 Cluster sampling

We cannot apply the covariate shift results to cluster sampling, because the data are not independent. Cluster sampling with unequal probabilities will require further research.

However, in the special case where the clusters themselves are sampled by SRS and the ultimate units are sampled by SRS within each cluster, then we can apply the methods of Dunn et al. (2022). Their paper is framed in terms of a more general two-layer hierarchical setting. They do not explicitly consider a finite-population setting, but their assumptions do allow for it (except for some restrictions in their CDF pooling method).

In the framework of Dunn et al. (2022), let $P_1, \dots, P_k \sim \Pi$ be k random distributions drawn iid from Π . From each of the sampled distributions P_ℓ for $\ell=1, \dots, k$, we draw n_ℓ iid observations $(X_{\ell 1}, Y_{\ell 1}), \dots, (X_{\ell n_\ell}, Y_{\ell n_\ell})$.

The corresponding setup in survey sampling would be cluster sampling, where our finite population of size N is partitioned into a fixed number K of clusters or Primary Sampling Units (PSUs), and we take a sample of these clusters. SRSWR from a finite population is a special case of iid sampling. Hence, if we take a SRSWR of $k < K$ clusters P_1, \dots, P_k from the finite population, and then take a SRSWR of n_ℓ ultimate units from cluster ℓ for each $\ell=1, \dots, k$, then this is a special case of the setup above, and we can apply most of the results in Dunn et al. (2022). We outline one of their approaches briefly here, and the others in Appendix B, but refer readers to their full paper for details.

Although Dunn et al. (2022) state their results in terms of iid sampling, it seems likely that they could be relaxed to exchangeable sampling of the distributions P_ℓ as well as exchangeable sampling within each

P_i . If so, these results would also apply to SRSWOR, not just SRSWR. Similarly, we conjecture that some of their results could be extended to the “weighted exchangeable” setting of Tibshirani et al. (2019).

Subsampling: By subsampling, Dunn et al. (2022) change the sampling design to become exchangeable. Start with the design above, but then subsample our dataset by choosing one unit at random from each cluster. Then any test case from any new cluster is exchangeable with our subsample. We can treat the k subsampled training cases and the one test case as being generated exchangeably by the process: “Take a cluster at random, then take one observation at random from that cluster”, and it is valid to use standard conformal methods.

However, although this process guarantees exact marginal coverage $1 - \alpha$ across training sets, it ignores most of the data and leads to wider variability in achieved coverage from training set to training set. An alternative is to carry out repeated subsampling B times and combine the results appropriately across subsamples. Dunn et al. (2022) show how to do this in a way that is guaranteed to have coverage of $1 - 2\alpha$, but in practice tends to achieve coverage close to $1 - \alpha$.

3.5 Stratified sampling

Again, we cannot apply the covariate shift results to stratified sampling, because even though strata are independent of each other, the n samples are not independent.

However, we can safely apply the conformal methods from previous subsections within each stratum separately, if within each stratum independently we have used one of the sampling methods with conformal guarantees. In other words, if the full population is partitioned into H strata, we can treat each stratum $h = 1, \dots, H$ as its own population. To form a prediction interval for a test case from stratum h , we apply conformal methods to only the n_h training cases from that stratum. Clearly this will guarantee conditional coverage by stratum. If the same coverage level is used simultaneously across all strata, it will also guarantee marginal coverage. This is an example of “group-balanced conformal prediction” or “object-conditional Mondrian conformal prediction” (Vovk, 2013; Vovk et al., 2022; Angelopoulos and Bates, 2022).

This stratum-by-stratum approach will lead to a loss of statistical efficiency, since each stratum’s conformal quantiles will be estimated using a sample size $n_h < n$. On the other hand, in some situations, prediction intervals may be more useful if we allow their sizes to vary by stratum than if their size has to be constant across strata. Further, guaranteeing coverage conditional on stratum may be more desirable than only guaranteeing marginal coverage, which could be achieved by overcoverage in some strata at the expense of undercoverage in others.

3.6 Post-stratification

Imagine our first n observations were drawn SRSWR, but we wish to post-stratify after data collection, and the population size N_h of each post-stratum is known. We could reweight each sampled observation

by the relative stratum sizes: unit-level post-stratification weights are proportional to N_h/n_h , where n_h is random rather than fixed in advance (Lohr, 2021).

We cannot use such weights for conformal inference and retain our exact finite-sample guarantees under the justifications in the present paper, because it would induce dependence between the training set and test case. However, Fannjiang et al. (2022) extend conformal methods to allow for “feedback covariate shift”, where the test distribution is allowed to depend on the observed training data, and this may be a promising direction for future work on post-stratified conformal prediction.

In the meantime, we can treat post-stratification as an approximation to estimating the covariate-shift likelihood ratio weights. Although we lose the exact conformal guarantees, using such an approximation would be just as reasonable as the estimation of covariate-shift weights in general. Specifically, imagine we are sampling SRSWR from two different finite populations: a training population of size M , and a test population of size N . Both populations have the same $P_{(X,Y)|S}$ and the same set of post-strata $1, \dots, H$, but different (and known) post-stratum sizes M_h, N_h for $h=1, \dots, H$. In both cases, our sampling design is equivalent to first choosing a post-stratum at random with probability proportional to post-stratum size, then a unit from within that post-stratum at random.

Now, we apply Lemma 3 – except that we let P_S and \tilde{P}_S depend on the post-stratum ID $h \in 1, \dots, H$, not the population unit ID. Then $dP_S/d\tilde{P}_S = (M_h/M)/(N_h/N)$, so $w(S_i) \propto N_h/M_h$. So far, we have exact guarantees. If we now assume that we do not actually know the true post-stratum sizes for the training population, we can replace M_h/M with training-sample estimates n_h/n and get post-stratification weights $w(S_i) \propto N_h/n_h$. If we further assume that both populations are actually the same, we are now justified in using conformal methods with the standard post-stratification weights. Our guarantees are approximate only because we estimated the “training” post-stratum sizes.

Similar types of weighting could also be developed in order to apply conformal inference when we do not wish to assume that test cases will be sampled uniformly but with some other sampling design. For instance, instead of guaranteeing prediction interval coverage across people, perhaps we want to guarantee coverage across visits to the doctor, and we use a sampling distribution to encode our knowledge of different people’s propensities to visit the doctor.

4. Examples

Our R code and knitted RMarkdown output are available at <https://github.com/ColbyStatSvyRsch/surveyConformal-paper-code>.

4.1 Real data

We have claimed that conformal methods may work better when they account for the sampling design of the data. As a simple demonstration, we turn to an extract of the Medical Expenditure Panel Survey or

MEPS (Agency for Healthcare Research and Quality, 2017), which is a nationally representative survey about the cost and use of health care among the U.S. civilian noninstitutionalized population.

We chose the MEPS because it has already been used as a benchmark dataset in several conformal inference papers, starting with Romano et al. (2019) and followed by others (Sesio and Candès, 2020; Sesio and Romano, 2021; Feldman et al., 2021; Bai et al., 2022). In each of these papers, the authors randomly partition MEPS data into proper training, calibration, and test sets, then report the coverage and length of conformal prediction intervals (PIs) for various models across many such random partitions. However, none of these papers report accounting for the complex sampling design of MEPS, which includes stratification, clustering, and oversampling of selected subgroups.

At present, we do not attempt a full correction of these earlier analyses of MEPS. We only wish to illustrate that there can be noticeable differences in the conformal PIs depending on whether or not we account for the sampling design, even in a very simple analysis. We use a portion of the public-use dataset for calendar year 2015. We subset to only those respondents who filled out the self-administered questionnaire (SAQ) portion of the survey, and we use the person-level weight variable designed to be used with the SAQ for persons age 18 and older during the interview.

In the poster associated with Romano et al. (2019), available at https://github.com/yromano/cqr/blob/master/poster/CQR_Poster.pdf, the authors explain that they are predicting “health care utilization, reflecting # visits to doctor’s office/hospital”. Following their GitHub code, we define a “utilization” response variable as the sum of five counts for 2015: total number of office-based visits reported; total number of reported visits to hospital outpatient departments; count of all emergency room visits reported; total number of nights associated with hospital discharges; and total number of days where home health care was received from any type of paid or unpaid caregiver.

Unlike the earlier conformal analyses of MEPS, we do take into account the public-use variables for strata, PSUs, and person-level weights. In the 2015 SAQ-eligible subset that we work with, there are 165 strata, and most have 2 or 3 PSUs. First we drop the 4 strata which had no observations in either PSU 1 or PSU 2. Next, for simplicity, we set aside every observation whose PSU is labeled 3 (regardless of stratum) and treat them as our test set. We treat the rest (PSUs 1 and 2) as our overall training set. We split this training data into proper-training and calibration sets under two different approaches. The first approach is a 50/50 SRS split that ignores the survey design. The second approach is to form a random split by PSU within each stratum, so that in each stratum independently, we randomly assign either PSU 1 to proper-training and PSU 2 to calibration or vice versa. These approaches result in proper-training and calibration sets with around 10,000 people each and a test set with 1,659 people.

For each split, we fit a linear regression model to the proper-training set to predict utilization from a subset of the covariates used by Romano et al. (2019): age; sex; indicators for diabetes diagnosis, private insurance coverage, and public insurance coverage; and quantitative summaries of answers to the Physical Component Summary (PCS), the Mental Component Summary (MCS), and the Kessler Index (K6) of

non-specific psychological distress. Higher PCS and MCS scores represent better health, while lower K6 scores represent less distress. We calculate PIs for each test case by combining it with the calibration set and finding conformal quantiles. Developing conformal methods for designs with only one or two PSUs per stratum is still an open problem, not yet addressed by the methods of Sections 3.4 and 3.5. Collapsing strata into pseudo-strata could be a reasonable solution if we had subject matter knowledge of the strata, but the public-use MEPS data uses anonymized stratum IDs. For this reason, our quantiles do not use these methods to handle the clustering and stratification, but they do apply the survey weights as in Section 3.2.

Because we ended up with such large proper-training and calibration sets, relative to the smaller test set, we only saw small differences between the exchangeable and design-based conformal approaches at moderate PI levels. On the other hand, we have enough proper-training and calibration data to estimate 99% PI levels too, and there we do see substantial differences in the average PI length. Our main takeaways, based on Tables 4.1 and 4.2:

1. When we used a conformal pipeline that assumed exchangeability (data splits at random; no weights in model-fitting; no weights in the conformal quantiles on the calibration set), we tended to over-cover. If we also ignored the weights when using the test set to estimate coverage, these calculations underestimated just how much over-coverage there was. By taking survey-weighted means on the test set, we found slightly higher estimates of coverage. We believe these higher estimates are more appropriate, since the survey-weighted means ought to generalize to the rest of the population better than un-weighted means do.
2. When we did use design-based methods (design-based splits; design-based and survey-weighted model fits; and survey-weighted conformal quantiles on the calibration set), this reduced our over-coverage a little, according to the survey-weighted means of coverage on the test set. Similarly, it also made our PIs a little narrower (around one to three fewer utilizations/year) for moderate PI levels, and substantially narrower (around forty fewer utilizations/year) for 99% PIs.

This brief MEPS example demonstrates that the estimated PI coverages and lengths can differ when conformal methods account for the survey design. In the next subsection, we study these effects in more detail, by repeatedly sampling under known sampling designs from a complete finite population. We also use smaller sample sizes, to see more pronounced differences between using vs. ignoring the survey design.

Table 4.1
Linear regression models' PI coverage, estimated on the MEPS dataset.

Split/fit/conformal	Test set	80% PI covg	90% PI covg	95% PI covg	99% PI covg
SRS	SRS	(0.824, 0.826)	(0.910, 0.911)	(0.951, 0.951)	(0.993, 0.994)
SRS	Svy-wtd	(0.836, 0.838)	(0.914, 0.915)	(0.950, 0.951)	(0.995, 0.995)
Svy-wtd	Svy-wtd	(0.829, 0.831)	(0.911, 0.913)	(0.949, 0.950)	(0.992, 0.993)

Notes: SRS = simple random sampling; PI = prediction interval; MEPS = Medical Expenditure Panel Survey. Coverages reported as approximate 95% CIs for the average, based on 100 random proper-training/calibration splits at each setting, using the same test set each time. When data splits, proper-training-set model fits, and calibration-set conformal quantiles ignored the survey design, we over-covered (especially for lower PI levels); but when test-set estimates of coverage also ignored the sampling design, they *underestimated* the amount of overcoverage, compared to test-set estimates that did account for the sampling design. However, when splits, fits, and conformal quantiles accounted for the survey design, there was slightly less over-coverage.

Table 4.2
Linear regression models' PI lengths, estimated on the MEPS dataset.

Split/fit/conformal	Test set	80% PI length	90% PI length	95% PI length	99% PI length
SRS	Either	(28.9, 29.2)	(43.9, 44.3)	(60.2, 60.7)	(250.5, 256.6)
Svy-wtd	Svy-wtd	(27.3, 27.7)	(40.9, 41.4)	(57.4, 58.0)	(202.0, 211.1)

Notes: SRS = simple random sampling; PI = prediction interval; MEPS = Medical Expenditure Panel Survey. PI lengths reported as approximate 95% CIs for the average, based on 100 random proper-training/calibration splits at each setting, using the same test set each time. When data splits, proper-training-set model fits, and calibration-set conformal quantiles ignored the survey design, our PI lengths tended to be slightly larger than when splits, fits, and conformal quantiles did account for the survey design – or much larger when the PI level is very high. In the Table's first row, it does not matter whether or not test-set estimates were survey-weighted, because these PI lengths are constant across test-set cases for a given data split and PI level.

4.2 Simulations

For our design-based simulations, we used the Academic Performance Index (API) data (California Department of Education, 2018) from R's `survey` package (Lumley, 2021). The `apipop` dataset contains information on 37 variables for all 6194 California schools (elementary, middle, or high school) with at least 100 students. The dataset vintage is not documented, but appears to be the 1999-2000 academic year, since the data includes API scores for each school for 1999 and 2000.

We used the `apipop` dataset as the finite population, and repeatedly took samples (with around $n = 200$ ultimate sampling units) using various designs. When we evaluated our results on test sets, we used the entire finite population – including those cases that had already been used to fit models or find conformal quantiles – because this corresponds to the guarantees that our paper makes in Section 3. Simulation details:

- In all simulations, sampling was done without replacement. Although our results in Section 3 assume sampling with replacement, we conjectured that sampling without replacement would still lead to conformal coverage close to nominal, and we wanted to check this empirically.
- All simulations were run 1,000 times. All 95% confidence intervals are calculated as the estimate ± 2 times the SD over $\sqrt{1,000}$.
- After dropping the rows with missing values for `enroll` and `mobility`, the full “finite population” consisted of the 6,153 schools without missing values for any variables used in the simulations.
- Most simulations used a sample size of $n = 200$ schools. However, the cluster samples had $n \approx 200$ schools on average but varied across samples. The regression model simulations took PPS samples of size $m + n = 400$, then split the samples at random into a proper training set of $m = 200$ and a calibration set of $n = 200$.
- The response variable was usually `api00`, the school's API in 2000. The exception is Table 4.5, where the response variable was `enroll`, the same variable used to create the PPS weights. For the non-regression simulations, we simply sought “unsupervised” prediction intervals for the marginal distribution of the response variable (with no covariates). For the regression simulations, we found quantiles of $|y - \hat{f}(x)|$ on the calibration set and sought prediction intervals for the response variable at the covariate values for each unit in the population.

- PPS sampling probabilities (if used) were usually proportional to `enroll`, the number of students enrolled at the school. The exception is parts of Tables 4.8 and 4.9, where PPS probabilities were proportional to 1 plus the square root of the residuals from the full-population linear regression model, in order to see the effects of over-sampling cases that are hard to fit well. Conformal quantiles for the PPS simulations were calculated as in Section 3.2.
- Clusters (if used) were based on `dnum`, the school district number. Cluster samples always took a SRS of 24 school districts. 24 was chosen because it led to an average of $n=198$ schools (close to the $n=200$ used in other sampling designs). For the “survey-design-aware” results in Table 4.6, we calculated quantiles using the “subsampling once” method as in Section 3.4, while the design-unaware results ignored clustering and calculated quantiles on the whole dataset.
- Strata (if used) were based on `stype`, the school type. Stratified samples always took 100 elementary, 50 middle, and 50 high schools, with an SRS within each school type. For the “survey-design-aware” results in Table 4.7, we calculated quantiles separately by stratum as in Section 3.5, while the design-unaware results ignored strata and calculated quantiles on the whole dataset.
- Linear regression models always predicted `api00` using a linear combination of `ell` (the percentage of English language learners), `meals` (the percentage of students eligible for subsidized meals), and `mobility` (the percentage of students for whom this is the first year at the school).

Our main takeaways:

1. Across many settings, using the naive quantile (the $\lceil n\alpha \rceil$ order statistic) instead of the conformal quantile (the $\lceil (n+1)\alpha \rceil$ order statistic) tended to give slight undercoverage. The conformal quantile helped partly to fix this; but in non-SRS settings it was not enough of a fix on its own.
2. For SRS designs, the conformal quantile lemma worked as advertised. See Table 4.3.
3. For PPS designs, ignoring the weights gave slight undercoverage when weights were not highly informative about the response variable. On the other hand, ignoring the weights led to extreme *over*-coverage when weights were highly informative. In both cases, weighted conformal quantiles fixed the problem. See Tables 4.4 and 4.5.
4. For clustered designs, as well as for stratified designs, ignoring the design undercovered but accounting for the design (including conformal-quantile padding) did fix it. The one exception was for one of the cluster-design simulations, where the design-based conformal PIs did not reach the target coverage. This might have been due to the small number of clusters, large variation in cluster sizes, and our choice of “subsampling once” as the conformal method. See Tables 4.6 and 4.7.
5. For simple models and split-conformal inference, if the weights were not highly informative about model variables or the fit of the model, then it did not make much difference whether or not the weights were used for quantiles. But when the weights were informative, we saw that unweighted conformal quantiles *over*-covered (and PIs were too wide). Using survey-weights in model-fitting was not enough to fix it, but weighting the quantiles was. See Tables 4.8 and 4.9.

Table 4.3

SRS. Average PI coverage of api00 , at two different PI levels, under 1,000 SRS samples of $n = 200$ each from API dataset.

Conformal?	80% PI coverage	90% PI coverage
no	(0.794, 0.801)	(0.945, 0.949)
yes	(0.800, 0.807)	(0.949, 0.953)

Notes: Coverages reported as 95% confidence intervals. Naive quantiles undercover, but conformal quantiles achieve target coverage.

Table 4.4

Uninformative PPS. Average PI coverage of api00 , at two different PI levels, under 1,000 PPS samples of $n = 200$ each from API dataset where probability $\propto \text{enroll}$.

Survey-weighted?	Conformal?	80% PI coverage	90% PI coverage
no	no	(0.752, 0.760)	(0.927, 0.932)
no	yes	(0.759, 0.767)	(0.935, 0.940)
yes	no	(0.795, 0.803)	(0.944, 0.949)
yes	yes	(0.803, 0.811)	(0.953, 0.958)

Notes: Coverages reported as 95% confidence intervals. Naive quantiles *under*-cover; conformal quantiles alone or survey-weighting alone do not fix it; but survey-weighted conformal quantiles achieve target coverage.

Table 4.5

Informative PPS. Average PI coverage of enroll , at two different PI levels, under 1,000 PPS samples of $n = 200$ each from API dataset where probability $\propto \text{enroll}$.

Survey-weighted?	Conformal?	80% PI coverage	90% PI coverage
no	no	(0.933, 0.935)	(0.986, 0.987)
no	yes	(0.934, 0.937)	(0.988, 0.989)
yes	no	(0.792, 0.798)	(0.946, 0.949)
yes	yes	(0.796, 0.802)	(0.948, 0.951)

Notes: Coverages reported as 95% confidence intervals. Naive quantiles *over*-cover; conformal quantiles alone do not fix it, while survey-weighting alone *under*-covers; but survey-weighted conformal quantiles achieve target coverage.

Table 4.6

Clustering. Average PI coverage of api00 , at two different PI levels, under 1,000 clustered samples of 24 clusters ($n \approx 200$) each from API dataset.

Survey-design?	Conformal?	80% PI coverage	90% PI coverage
no	no	(0.785, 0.805)	(0.934, 0.943)
no	yes	(0.791, 0.811)	(0.940, 0.949)
yes	no	(0.799, 0.817)	(0.916, 0.930)
yes	yes	(0.799, 0.817)	(0.959, 0.968)

Notes: Coverages reported as 95% confidence intervals. Due to high variability in cluster sizes, these 95% CIs are wider than in previous tables, but overall trend is generally similar to other tables: Naive quantiles appear likely to *under*-cover; conformal quantiles alone or survey-design-aware analyses alone do not necessarily fix it; but survey-design-aware conformal quantiles achieve target coverage.

Table 4.7

Stratification. Average PI coverage of api00 , at two different PI levels, under 1,000 stratified samples of $n = 200$ each from API dataset.

Survey-design?	Conformal?	80% PI coverage	90% PI coverage
no	no	(0.769, 0.777)	(0.933, 0.938)
no	yes	(0.775, 0.783)	(0.940, 0.944)
yes	no	(0.787, 0.795)	(0.939, 0.944)
yes	yes	(0.800, 0.808)	(0.953, 0.956)

Notes: Coverages reported as 95% confidence intervals. Naive quantiles *under*-cover; conformal quantiles alone or survey-design-aware analyses alone do not fix it; but survey-design-aware conformal quantiles achieve target coverage.

Table 4.8

Linear regression models' PI coverage. Average PI coverage of `api00`, at two different PI levels, under 1,000 PPS samples of $n = 200$ each from API dataset.

PPS probs	Svy-wtd conformal?	Svy-wtd regression?	80% PI covg	90% PI covg
enroll	no	no	(0.807, 0.810)	(0.955, 0.957)
enroll	yes	no	(0.803, 0.807)	(0.954, 0.956)
residuals	no	no	(0.874, 0.876)	(0.973, 0.974)
residuals	no	yes	(0.875, 0.878)	(0.973, 0.974)
residuals	yes	no	(0.798, 0.802)	(0.950, 0.951)
residuals	yes	yes	(0.798, 0.801)	(0.950, 0.951)

Notes: Coverages reported as 95% confidence intervals. For weights proportional to `enroll` (uninformative for the regression), it makes little difference whether or not we weight the conformal quantiles. For informative weights proportional to full-pop residuals, un-weighted conformal quantiles *over-cover*, and this is fixed by survey-weighted conformal quantiles; but it makes little difference whether or not we fit survey-weighted regression models.

Table 4.9

Linear regression models' PI lengths. Average length of PIs for `api00`, at two different PI levels, under 1,000 PPS samples of $n = 200$ each from API dataset.

PPS probs	Svy-wtd conformal?	Svy-wtd regression?	80% PI length	90% PI length
enroll	no	no	(192.6, 194.1)	(294.5, 297.3)
enroll	yes	no	(190.8, 192.7)	(293.7, 297.2)
residuals	no	no	(214.1, 215.6)	(328.2, 331.5)
residuals	no	yes	(211.8, 213.3)	(331.4, 334.8)
residuals	yes	no	(179.4, 180.8)	(285.4, 287.7)
residuals	yes	yes	(175.1, 176.3)	(286.1, 288.3)

Notes: Lengths reported as 95% confidence intervals. For weights proportional to `enroll` (uninformative for the regression), it makes little difference whether or not we weight the conformal quantiles. For informative weights proportional to full-pop residuals, PIs from un-weighted conformal quantiles are much wider than PIs from survey-weighted conformal quantiles; but it makes little difference whether or not we fit survey-weighted regression models.

Overall, the survey-conformal quantiles we proposed mathematically in Section 3 also appear to work empirically. Data analysts will likely get coverage closer to nominal when they account for the weights or other survey design features.

5. Extensions

We discuss several practical considerations: What if the sampling design does not quite match the situations above? What if the sampling probabilities are not all known? We also suggest other possible use cases for conformal inference in survey methodology.

5.1 Practical considerations

Weighting adjustments: Most surveys are not released with inverse-probability sampling weights alone. The final survey weights have been adjusted for nonresponse, post-stratification, and other considerations. Tibshirani et al. (2019) found that their conformal methods still maintained coverage close to nominal even when they only estimated the likelihood ratio weights, instead of using true likelihood ratios. We anticipate

similar outcomes for conformal methods that use adjusted sampling weights instead of the true inverse-probability sampling weights.

Distribution shift: Surveys are typically assumed to be sampled from a well-defined finite population, in a specific time and place. We may not be guaranteed coverage if we make predictions for units at future times or from distinct populations. If model-based (rather than design-based) inference makes more sense for the application at hand, we may still be able to estimate the covariate shift likelihood ratio and apply the conformal techniques of Tibshirani et al. (2019); it will simply no longer be strictly design-based inference, though a joint superpopulation/design-based framework may be a fruitful topic for future work (Isaki and Fuller, 1982; Rubin-Bleuer and Şchiopu-Kratina, 2005; Han and Wellner, 2021).

Unknown sampling probabilities for test cases: Our conformal methods of Section 3.2 require the sampling probabilities for each test case. However, in practice the sampling probabilities are not typically known for every population unit. Even if they are known internally within the survey organization, public release of the entire population's sampling probabilities can increase the risk of unit re-identification and privacy breaches.

Instead, the survey organization could report a set of population categories for which the sampling probabilities are approximately equal within each category, along with their approximate probabilities. There may be categories that are broad enough to minimize privacy risks, but fine enough to approximate the real sampling probabilities well. Or instead of discrete categories, the survey organization could report a kind of generalized variance function or GVF (Wolter, 2007), but used to estimate each population unit's sampling probabilities rather than variances, based on covariates available for each unit. This may be especially reasonable for PPS.

Alternatively, the survey organization could release only the value of the single largest inverse-probability sampling weight for the whole population. Then, using that one weight for every test case during conformal prediction would be strictly conservative. Or as an approximation, users could use the largest sampling weight in the public-use dataset. Users could also try a range of plausible weights for a desired test case, based on the weights of similar in-sample units, and report a sensitivity analysis.

Finally, if a user is interested in predictions for a population unit with the same vector of covariates X as one of the sampled cases in the dataset, they could simply use that sampled case's survey weight.

5.2 Other uses for conformal methods in survey methodology

National statistical offices, polling agencies, and any others who collect and pre-process survey data may find their own use cases for conformal methods:

- Build a response propensity model based on internal metadata from past surveys. For the next survey, build conformal prediction sets for each sampled unit's most likely mode of response (or ultimate status of nonresponse). Use these sets to help choose the mode of initial contact for each sampled unit.

- As part of automated quality control checks, use conformal prediction intervals or sets for each variable to flag potential outliers for followup (Bates, Candès, Lei, Romano and Sesia, 2023).
- To impute for item nonresponse, or to generate synthetic microdata, draw at random from a conformal prediction interval or set for that variable.

6. Conclusion

There is growing interest in extending machine learning (ML) methods to complex survey data, as well as a distinct body of work around developing prediction intervals for new predictive methods such as those arising from the ML community. While this ongoing work is valuable, design-based conformal inference provides an alternative approach for both needs: We can apply a novel predictive algorithm to complex survey data (even if that algorithm has not been specifically adapted to account for the survey design yet), and automatically get conformal intervals compatible with the fitted prediction function (even if native prediction-interval methods have not been developed for that algorithm yet), and still manage to provide exact, finite-sample, design-based coverage guarantees.

Of course, design-based conformal methods are not a panacea. First, for methods that do have well-understood design-based adaptations, the design-based version fitted to complex-survey-design training data is likely to be a better predictive model for the conditional mean and consequently to have narrower conformal prediction intervals than when that method ignores the sampling design. As a simple example, in certain cases a survey-weighted linear model may generalize better than an unweighted linear model fit to the same data, so the survey-weighted model's smaller residuals will lead to narrower conformal prediction intervals than for the unweighted model. In this sense, it is still important to develop survey-weighted equivalents of novel ML algorithms.

Second, for models that do have native prediction intervals, if we can justifiably trust that the underlying model assumptions are met, then we may be able to get narrower or less-variable prediction intervals natively than from conformal inference. And in some cases, native methods can give us conditional prediction intervals rather than marginal ones. In this sense, it is still important to develop native prediction intervals for specific models.

But when either of these conditions is not yet met, conformal inference is a practical and assumption-lean way to fill in the gap. We also note that Lei et al. (2018) found in simulations that conventional methods are so noisy for high-dimensional regression that conformal prediction intervals can actually be narrower than native ones.

We have also discussed several pragmatic limitations to conformal inference with survey data: Sampling weights are not truly known in advance, nor can they be reported for every population unit. Our suggested solutions are only a starting point.

Still, we hope that this paper spurs interest in conformal inference among survey statisticians. Survey data analysts might find these methods directly applicable. Furthermore, research into conformal methods

from the design-based perspective could bring new insights back to the wider statistics/ML community of conformal inference researchers.

As one example, Vovk (2013) shows that the distribution of achieved coverage levels has a Beta distribution across calibration sets when data are exchangeable. In order to gauge how much variability to expect from this Beta distribution, Angelopoulos and Bates (2022) suggest one estimate of the effective sample size n_{eff} for weighted conformal methods, but their estimate is not appropriate for all sampling designs nor for all estimators. Survey statisticians may be able to suggest better estimates of n_{eff} or recommend other ways to study and control the variability in achieved coverage for non-exchangeable data.

Other open problems include conformal methods for unequal-probability sampling WOR; unequal-probability cluster sampling; strata with few PSUs; combining strata rather than analyzing each stratum separately; exact coverage guarantees for post-stratification; panel survey designs; or joint superpopulation/design-based frameworks. We encourage survey researchers to contribute to conformal methodology and find new areas of application for these methods.

Acknowledgements

The author thanks the Editor, the Associate Editor, and the anonymous referees for their helpful suggestions which improved this manuscript. The author is also grateful to Ryan Tibshirani, Robin Dunn, Benjamin LeRoy, and Evan Randles for their early feedback on several of the ideas presented here.

Appendix

A. Adaptive prediction regions

In the simple approach to conformal prediction for regression described in our paper, we have used a constant-width prediction band everywhere, which might be unrealistic for many scenarios. If the true conditional distribution of $Y|X$ is heteroscedastic for instance, marginal coverage will still be correct, but it will be achieved by overcovering at some regions of X and undercovering at others. For “adaptive” alternatives, where the PI is wider at regions of X with more variability in Y , see recent work on locally-weighted conformal inference (Lei et al., 2018) and conformalized quantile regression (Romano et al., 2019).

For classification problems, even the simple approach of Section 2.5 produces adaptive prediction sets: for hard test cases where \hat{f} is uncertain about the right class, prediction sets will be larger than for easy test cases where \hat{f} confidently assigns most of the probability to one class. See also Romano et al. (2020) and Angelopoulos and Bates (2022).

Varying PI widths due to survey weighting: Note that in the survey-weighted setting, we will automatically get slightly different prediction interval widths at different test cases, because the survey weights can differ for each unit being predicted. However, the effect of survey weighting is to widen slightly the

prediction intervals for regions of X with a smaller effective sample size. This is distinct from adaptivity to regions with more variability in Y .

Loosely, a sampled unit with a larger sampling weight represents more population units, so we are more uncertain about units “like” this one. When we apply Lemma 3 to a test case with a larger weight, its survey-weighted eCDF is pushed down farther than if it had a small weight; so the nonconformity score quantile is estimated farther to the right; so the conformal prediction interval is wider for units with larger weights, all else being equal.

On the other hand, consider optimal allocation designs, in which strata with higher variance of Y are oversampled. When we construct conformal intervals separately within each stratum, the high-variance strata will have their eCDFs padded by a smaller $1/n_h$ than low-variance strata do, so the quantile adjustment is less conservative. But this should be more than offset by the fact that high-variance strata also have wider spread in Y , so that ultimately their prediction intervals will end up wider than those of low-variance strata.

B. Other conformal methods for cluster samples

Beyond Section 3.4, we briefly note how Dunn et al. (2022)’s other methods relate to cluster designs in survey sampling.

Prediction for an observed cluster: If we only need a prediction interval for new observations from one of the clusters we already sampled, we can simply apply standard conformal methods by only using that cluster’s data, which will be exchangeable.

Double conformal: For unsupervised prediction (where we want a prediction interval for Y without conditioning on covariates X) for a new cluster, one could create conformal prediction intervals separately within each cluster, then combine their endpoints appropriately into a single interval. This lines up with the design-based spirit: construct valid estimates within each cluster, then combine them sensibly across clusters. Dunn et al. (2022) derive such a method that is guaranteed to have coverage at least $1-\alpha$. However, in simulations it overcovers, with coverage of nearly 1 and wider intervals than other approaches.

Pooling CDFs: We could first construct eCDFs within each cluster and average them together into one pooled eCDF. Then we could apply standard conformal methods using this pooled eCDF. Dunn et al. (2022) only prove that this is asymptotically valid, requiring a continuous distribution for Y as well as a growing number of sampled clusters $k \rightarrow \infty$. This is not possible in the traditional design-based setting of a fixed finite population, although we could construct a superpopulation model and a sequence of growing finite populations that satisfies their requirements.

In simulations, Dunn et al. (2022) find that CDF pooling tends to have coverage closest to nominal as well as shortest length of prediction intervals across most settings. But if we do not wish to rely on

asymptotic arguments and continuous Y data, repeated subsampling appears to work better than single subsampling or the double conformal method.

References

- Agency for Healthcare Research and Quality (2017). Medical Expenditure Panel Survey 2015 full year consolidated data file. https://meps.ahrq.gov/mepsweb/data_stats/download_data_files_detail.jsp?cboPufNumber=HC-181.
- Angelopoulos, A.N., and Bates, S. (2022). A gentle introduction to conformal prediction and distribution-free uncertainty quantification. *arXiv preprint arXiv:2107.07511*.
- Bai, Y., Mei, S., Wang, H., Zhou, Y. and Xiong, C. (2022). Efficient and differentiable conformal prediction with general function classes. *arXiv preprint arXiv:2202.11091*.
- Barber, R.F., Candès, E.J., Ramdas, A. and Tibshirani, R.J. (2021). The limits of distribution-free conditional predictive inference. *Information and Inference: A Journal of the IMA*, 10(2), 455-482.
- Barber, R.F., Candès, E.J., Ramdas, A. and Tibshirani, R.J. (2022). Conformal prediction beyond exchangeability. *arXiv preprint arXiv:2202.13415*.
- Bates, S., Candès, E., Lei, L., Romano, Y. and Sesia, M. (2023). Testing for outliers with conformal p-values. *The Annals of Statistics*, 51(1), 149-178.
- Bechtold, W.A., and Patterson, P.L. (2005). *The Enhanced Forest Inventory and Analysis Program – National Sampling Design and Estimation Procedures*.
- Bersson, E., and Hoff, P.D. (2022). Optimal conformal prediction for small areas. *arXiv preprint arXiv:2204.08122*.
- California Department of Education (2018). Academic Performance Index. <https://www.cde.ca.gov/re/pr/api.asp>.
- Central Statistical Office, Zambia (2015). Living Conditions Monitoring Survey (LCMS). <https://www.zamstats.gov.zm/portfolio/living-conditions-monitoring-survey-lcms/>.
- Cherian, J., and Bronner, L. (2020). How the Washington Post estimates outstanding votes for the 2020 presidential election. https://s3.us-east-1.amazonaws.com/elex-models-prod/2020-general/write-up/election_model_writeup.pdf.
- Chernozhukov, V., Wüthrich, K. and Yinchu, Z. (2018). Exact and robust conformal inference methods for predictive machine learning with dependent data. *Conference On Learning Theory*, 732-749. PMLR.

- Crookston, N.L., and Dixon, G.E. (2005). The forest vegetation simulator: A review of its structure, content, and applications. *Computers and Electronics in Agriculture*, 49(1), 60-80.
- Dagdoug, M., Goga, C. and Haziza, D. (2021). Model-assisted estimation through random forests in finite population sampling. *Journal of the American Statistical Association*.
- Dagdoug, M., Goga, C. and Haziza, D. (2022). Model-assisted estimation in high-dimensional settings for survey data. *Journal of Applied Statistics*.
- Dunn, R., Wasserman, L. and Ramdas, A. (2022). Distribution-free prediction sets for two-layer hierarchical models. *Journal of the American Statistical Association*.
- Durrett, R. (2019). *Probability: Theory and Examples*. Cambridge University Press.
- Fannjiang, C., Bates, S., Angelopoulos, A.N., Listgarten, J. and Jordan, M.I. (2022). Conformal prediction under feedback covariate shift for biomolecular design. *PNAS*, 119(43), e2204569119.
- Feldman, S., Bates, S. and Romano, Y. (2021). Improving conditional coverage via orthogonal quantile regression. *arXiv preprint arXiv:2106.00394*.
- Fienberg, S.E. (2010). The relevance or irrelevance of weights for confidentiality and statistical analyses. *Journal of Privacy and Confidentiality*, 1(2).
- Fong, E., and Holmes, C.C. (2021). Conformal Bayesian computation. *Advances in Neural Information Processing Systems*, 34, 18268-18279.
- Gibbs, I., and Candes, E. (2021). Adaptive conformal inference under distribution shift. *Advances in Neural Information Processing Systems*, 34, 1660-1672.
- Han, Q., and Wellner, J.A. (2021). Complex sampling designs: Uniform limit theorems and applications. *The Annals of Statistics*, 49(1), 459-485.
- Hong, H.G., and He, X. (2010). Prediction of functional status for the elderly based on a new ordinal regression model. *Journal of the American Statistical Association*, 105(491), 930-941.
- Isaki, C.T., and Fuller, W.A. (1982). Survey design under the regression superpopulation model. *Journal of the American Statistical Association*, 77(377), 89-96.

- Krebs, M.A., Reeves, M.C. and Baggett, L.S. (2019). Predicting understory vegetation structure in selected western forests of the United States using FIA inventory data. *Forest Ecology and Management*, 448, 509-527.
- Kshirsagar, V., Wiecek, J., Ramanathan, S. and Wells, R. (2017). Household poverty classification in data-scarce environments: A machine learning approach. *arXiv preprint arXiv:1711.06813*.
- Lei, J., G'Sell, M., Rinaldo, A., Tibshirani, R.J. and Wasserman, L. (2018). Distribution-free predictive inference for regression. *Journal of the American Statistical Association*, 113(523), 1094-1111.
- LeRoy, B., and Schafer, C. (2021). Conformal prediction for simulation models. *ICML 2021 Workshop on Distribution-Free Uncertainty Quantification*. https://benjaminleroy.github.io/documents/icml2021/conformal_prediction_for_simulation_models.pdf.
- Lohr, S.L. (2021). *Sampling: Design and Analysis*. Chapman and Hall/CRC.
- Lumley, T. (2021). survey: Analysis of complex survey samples. R package version 4.1-1.
- Lumley, T., and Scott, A. (2017). Fitting regression models to survey data. *Statistical Science*, 265-278.
- Lunde, R. (2023). On the validity of conformal prediction for network data under non-uniform sampling. *arXiv preprint arXiv:2306.07252*.
- Manokhin, V. (2022). Awesome conformal prediction. <https://github.com/valeman/awesome-conformal-prediction>.
- McConville, K.S., Breidt, F.J., Lee, T. and Moisen, G.G. (2017). Model-assisted survey regression estimation with the lasso. *Journal of Survey Statistics and Methodology*, 5(2), 131-158.
- National Center for Health Statistics (2016). *The Second Longitudinal Study of Aging (LSOA II)*. <https://www.cdc.gov/nchs/lsoa/lsoa2.htm>.
- Oliveira, R.I., Orenstein, P., Ramos, T. and Romano, J.V. (2022). Split conformal prediction for dependent data. *arXiv preprint arXiv:2203.15885*.
- Papadopoulos, H., Proedrou, K., Vovk, V. and Gammerman, A. (2002). Inductive confidence machines for regression. *European Conference on Machine Learning*, 345-356. Springer.

- Romano, Y., Patterson, E. and Candes, E. (2019). Conformalized quantile regression. *Advances in Neural Information Processing Systems*, 32.
- Romano, Y., Sesia, M. and Candes, E.J. (2020). Classification with valid and adaptive coverage. *arXiv preprint arXiv:2006.02544*.
- Rubin-Bleuer, S., and Şchiopu-Kratina, I. (2005). On the two-phase framework for joint model and design-based inference. *The Annals of Statistics*, 33(6), 2789-2810.
- Sande, L.S., and Zhang, L.-C. (2021). Design-unbiased statistical learning in survey sampling. *Sankhyā A*, 83(2), 714-744.
- Särndal, C.-E., Swensson, B. and Wretman, J. (1992). *Model Assisted Survey Sampling*. Springer.
- Sesio, M., and Candes, E.J. (2020). A comparison of some conformal quantile regression methods. *Stat.*, 9(1), e261.
- Sesio, M., and Romano, Y. (2021). Conformal prediction using conditional histograms. *Advances in Neural Information Processing Systems*, 34.
- Thompson, S.K. (1997). Adaptive sampling in behavioral surveys. In *The Validity of Self-Reported Drug Use: Improving the Accuracy of Survey Estimates*, volume 167 of *NIDA Research Monograph*, (Eds., L. Harrison and A. Hughes), 296-319. National Institutes of Health.
- Tian, Q., Nordman, D.J. and Meeker, W.Q. (2022). Methods to compute prediction intervals: A review and new results. *Statistical Science*, 37(4), 580-597.
- Tibshirani, R.J., Barber, R.F., Candes, E. and Ramdas, A. (2019). Conformal prediction under covariate shift. *Advances in Neural Information Processing Systems*, 32.
- Tu, S. (2014). The Dirichlet-Multinomial and Dirichlet-Categorical models for Bayesian inference. <https://stephentu.github.io/writeups/dirichlet-conjugate-prior.pdf>. Unpublished manuscript.
- Vovk, V. (2013). Conditional validity of inductive conformal predictors. *Machine Learning*, 92, 349-376.
- Vovk, V., Gammerman, A. and Shafer, G. (2022). *Algorithmic Learning in a Random World*. Springer, 2nd edition.

Wieczorek, J., Guerin, C. and McMahon, T. (2022). K-fold cross-validation for complex sample surveys. *Stat*, 11(1), e454.

Wolter, K. (2007). Generalized variance functions. *Introduction to Variance Estimation*. Springer, 2nd edition.

Yang, S., and Kim, J.K. (2019). Nearest neighbor imputation for general parameter estimation in survey sampling. *The Econometrics of Complex Survey Data*, volume 39 of *Advances in Econometrics*, 209-234. Emerald Publishing Limited.